

TECHNIQUE FOR MULTIPROTOCOL TRANSPORT USING MPLS

(MULTI-PROTOCOL LABEL SWITCHING)

Field of the Invention

5 The present invention relates to a method and/or
architecture for network routing generally and, more particularly,
to a router and method for sending multiple protocols over a single
pipe within the network.

10 Background of the Invention

Traditional Internet Protocol (IP) based networks require
every node to look at network layer values of an Open Systems
Interconnection (OSI) model (i.e., IP addresses) within each frame,
refer to local network routing databases, and then make decisions
15 about where to forward the frames. With a large number of nodes
and networks, management of the local network routing tables
becomes a difficult task, both for hardware and software. In
addition, current IP networks provide a single routed path between
any source and a destination. Therefore, even if network bandwidth

0325.00526
CD01223

is available, current protocols are not able to utilize different paths through the IP network efficiently.

Referring to FIG. 1, a drawing of an Ethernet frame 10 routed with a conventional Multi-Protocol Label Switching (MPLS) protocol is shown. The MPLS protocol has been developed for solving the routing table and single path problems. The MPLS protocol inserts a 16-bit protocol identification field 12 and one or more 32-bit shim headers 14 between an OSI data layer (i.e., layer 2) address 16 and an OSI network layer (i.e., layer 3) address 18. Each of the shim headers 14 contains a 20-bit value called an MPLS label 20. An OSI network layer (i.e., layer 3) protocol identification field 22 of the original Ethernet frame 10 is discarded in the process.

Referring to FIG. 2, a drawing of a Point-to-Point Protocol (PPP) frame 24 routed with the conventional MPLS protocol is shown. The PPP frame 24 is processed for the MPLS protocol by removing an OSI network layer protocol identification field 26. The MPLS protocol identification field 12 and one or more shim headers 14 are then added.

A path formed by the MPLS labels 20 is called a Label Switched Path (LSP). Signaling protocols establish paths and

0325.00526
CD01223

assign the 20-bit values at nodes along the paths. In addition, signaling protocols such as Resource Reservation Protocol (RSVP) have been extended to allow path reservation using MPLS to support proper Traffic Engineering (TE) over the paths. Using the header values to identify routes, network nodes are able to avoid processing the network layer (i.e., IP) addresses at every node to determine the path for a frame.

Currently, MPLS is being used for edge, core, and long-haul networks - for both legacy as well as optical networks. Carriers are using MPLS as a foundation for next-generation network offerings. Versions of MPLS are also targeted for replacing the control plane for Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) and optical cross-connects. MPLS switching, MPLS-based data transport, MPLS virtual private networks, and so on, are presently big business. Many companies are involved in designing Complex Programmable Logic Devices (CPLDs), Application Specific Integrated Circuits (ASICs), Field Programmable Gate Arrays (FPGAs), boards, and systems based on MPLS.

A problem, however, is limited support for multiprotocol transport in MPLS. The MPLS protocol predominantly carries IP traffic. Other protocols can be sent, but to do so requires

0325.00526
CD01223

changes to signaling protocols. Even when the signal protocols are changed, only a single type of protocol can be carried over a single path. Each time that a new protocol is added, the signaling software needs to be upgraded at every node to account for the new protocol.

In conventional protocols, the moment a router at the edge of an MPLS network adds the MPLS labels the original OSI network layer protocol identification 22/26 is lost. Thereafter, the only way for a terminal router at the end of the MPLS network to determine the OSI network layer protocol identification 22/26 required to reconstruct the Ethernet frame 10 or PPP frame 24 is to rely on signaling mechanisms within the MPLS protocol. The MPLS signaling mechanisms negotiate a protocol to be transported on an LSP. Therefore, different LSPs are required for transporting different protocols. In addition, every node along the path needs to be aware of any new protocol being established, thus resulting in a large and dynamic software/hardware infrastructure.

Another problem arises with signaling protocols such as RSVP where a path needs to be periodically refreshed to keep the path alive. Because an individual LSP is required for each flow, many LSPs are commonly required between two end points. With many

0325.00526

CD01223

LSPs, a large number of refreshes cause heavy loads on processors and on the network. The above-mentioned problems have restricted wide adoption of MPLS. In addition, MPLS systems have been complex and expensive due to additional hardware complexity required for multi-service transport.

Summary of the Invention

The present invention concerns a router generally comprising a first port, a second port, and a circuit. The first port may be configured to receive a frame having a network layer protocol identification. The second port may be connectable to a Multi-Protocol Label Switching (MPLS) network. The circuit may be configured to (i) insert an MPLS label into the frame while retaining the network layer protocol identification and (ii) present the frame in the MPLS network per the MPLS label.

The objects, features and advantages of the present invention include providing a router and method that may provide for (i) sending multiple protocols/flows simultaneously in a Multi-Protocol Label Switching (MPLS) Label Switched Path, (ii) constructing paths for certain Traffic Engineering parameters over a network and/or multiple networks then using each path for

multiple types of traffic, (iii) simpler, less expensive, and more scalable hardware and software solutions for devices and networking systems, (iv) consuming fewer network resources to accommodate multi-protocol transport, (v) a capability where network providers can establish network paths for customers independent of the type of data the customers are sending, and/or (vi) no additional hardware and/or software burden on intermediate MPLS nodes for each end system protocol that is to be sent on MPLS paths.

Brief Description of the Drawings

These and other objects, features and advantages of the present invention will be apparent from the following detailed description and the appended claims and drawings in which:

FIG. 1 is a drawing of a conventional Ethernet frame in MPLS;

FIG. 2 is a drawing of a conventional PPP frame in MPLS;

FIG. 3 is a drawing of an Ethernet frame in MPLS in accordance with the present invention;

FIG. 4 is a drawing of a PPP frame in MPLS in accordance with the present invention;

FIG. 5 is a block diagram of a network;

FIG. 6 is a flow diagram of a method for inserting a frame into the network;

FIG. 7 is a flow diagram of a method for extracting the frame from the network; and

5 FIG. 8 is a block diagram of an example circuit implementing the insertion and extraction methods.

Detailed Description of the Preferred Embodiments

Referring to FIG. 3, a drawing is shown of an Ethernet frame in a Multi-Protocol Label Switching (MPLS) protocol 100 in accordance with a preferred embodiment of the present invention. The MPLS protocol is generally defined in the document, "Multiprotocol Label Switching Architecture", Internet Engineering Task Force (IETF), Reston Virginia, Request For Comment (RFC) 3031, 10
15 hereby incorporated by reference in its entirety. The MPLS protocol is generally modified by the present invention.

Modifying the Ethernet frame for the MPLS generally comprises adding a field 102 and a stack 104 to the Ethernet frame. The field 102 may be implemented as an MPLS protocol identification field as defined by the RFC 3031. The stack 104 may be implemented 20 as an MPLS label stack as defined by the RFC 3031.

The MPLS protocol identification field 102 and the MPLS label stack 104 may be inserted into the Ethernet frame between a field 106 and a field 108. The field 106 may contain an Open Systems Interconnection (OSI) model data link layer (e.g., layer 2) address. The field 108 may contain an OSI network layer (e.g., layer 3) address. The present invention may depart from the RFC 3031 by retaining a field 110 between the field 106 and the field 108. The field 110 may contain an OSI network layer protocol identification for the Ethernet frame. Therefore, MPLS switching may be made completely independent of the layer 3 protocol identification and other fields.

The MPLS label stack 104 may contain one or more headers 112. Each header 112 may comprise a label 114, a class of service 116, a flag (e.g., S) 118 for indicating a bottom of stack, and a Time To Live (TTL) value 120. The labels 114 of each header 112 may be used with a Label Switched Path (LSP) through an MPLS network (FIG. 5).

Referring to FIG. 4, a drawing is shown of a Point-to-Point Protocol (PPP) frame in MPLS 124. The PPP frame may be modified by adding the MPLS protocol identification field 102 and the MPLS stack 104 between a field 126 and a field 128. The field

126 may contain an OSI data link layer address. The field 128 may contain an OSI network layer address. A field 130 may be retained as the MPLS protocol identification field 102 and the MPLS stack fields 104 are added. The field 130 may contain an OSI network layer protocol identification for the PPP frame.

Referring to FIG. 5, a block diagram of a network 132 is shown. The network 132 generally comprises a physical layer 134 and two or more routers 136A-B. Additional nodes or routers 136 (not shown) may be provided within the network 132 between the router 136A and the router 136B. The network 132 may be implemented as an MPLS network. The router 136A may be implemented as an edge router of the MPLS network 132. The router 136B may be implemented as another edge router of the MPLS network 132.

The edge router 136A may have a port 137A for connecting to the MPLS physical layer 134. The edge router 136A may have another port 138A for connecting to another network 139. The network 139 may comprise one or more nodes 140A-B. Each node 140A-B may communicate in the network 139 with a similar and/or different protocol. The edge router 136B may have a port 137B for connecting to the MPLS physical layer 134. The edge router 136B may have another port 138B for connecting to another network 141.

0325.00526

CD01223

The network 141 may comprise one or more nodes 142A-B. Each node 142A-B may communicate in the network 141 with a similar and/or different protocol.

The MPLS network 132 may allow a node (e.g., node 140A) on the network 139 to communicate with another node (e.g., node 142A) on the network 141. Communications through the MPLS network 132 may be enabled by creating LSPs 144A-C through conventional signaling protocols. Signaling protocols such as Resource Reservation Protocol (RSVP) generally allow path reservations using MPLS to support proper Traffic Engineering (TE) over the LSPs 144A-C. Using the MPLS labels 112 to identify a particular LSP 144, the MPLS network nodes 136 may be able to avoid processing the OSI network layer addresses at every node 136 to determine the path for a frame.

In an example, an AppleTalk® frame originating from the node 140B may be sent to the edge router 136A over the network 139. The edge router 136A may operate as an ingress node to the MPLS network 132 for the AppleTalk® frame. A signaling protocol may establish the LSP 144C as the proper traffic-engineered path for the AppleTalk® frame to the edge router 136B. The edge router 136A may incorporate an appropriate MPLS protocol identification and an

0325.00526
CD01223

MPLS stack into the AppleTalk® frame. The AppleTalk® frame in MPLS may then be transferred along the LSP 144C per the MPLS labels to the edge router 136B. The edge router 136b may operate as an egress node from the MPLS network 132 for the AppleTalk® frame.

5 The edge router 136B generally remove the MPLS protocol identification and the MPLS stack. Since the layer 3 protocol identification information for the AppleTalk® frame has been retained, the edge router 136B may not need to recreate the layer 3 protocol identification from MPLS label value lookups. Thereafter, the edge router 136B may present the AppleTalk® frame on the network 141 where it may be received by the node 142B.

The operations of the edge router 136A and the edge router 136B in routing the AppleTalk® frame along the LSP 144C are generally independent of the protocol of the AppleTalk® frame. As
15 a result, the edge routers 136A and 136B may also route frames having other protocols along the same LSP 144C. The transfers along the LSP 144C may be unidirectional or bidirectional. For example, an IP frame originating from node 142A may be received by the edge router 136B. The edge router 136B, operating as an
20 ingress node, may insert the appropriate MPLS protocol identification and the MPLS stack into the IP frame. The IP frame

0325.00526
CD01223

in MPLS may then be transferred along the LSP 144C to the edge
router 136A. The edge router 136A, operating as an egress node,
may remove the MPLS protocol identification and the MPLS stack
without the need to recreate the layer 3 protocol identification
5 for the IP frame from MPLS label value lookups. The IP frame may
then be presented to the network 139 for reception by the node
140A. Likewise, an Internetwork Packet Exchange (IPX) frame may be
routed from the node 140A through the LSP 144C to the node 142A.

A result of the present invention may be that the MPLS
network 132 may carry the AppleTalk® frame, the IP frame, and the
IPX frame over the same LSP 144C. Traffic Engineering (TE)
parameters over a single MPLS network or multiple networks may be
use to establish LSPs 144 for sending multiple types of traffic.
The result is generally a simpler, cheaper, and more scalable
15 hardware and software solution for both devices and networking
systems. For example, network providers may give a traffic-
engineered path to a customer and let the customer use the path for
any type of data at any time. The network providers may not be
concerned with signaling protocols for any of the protocols used by
20 the customer. The path may now be used by the customer as a pipe
for any purpose without additional burdens on the network provider.

Referring to FIG. 6, a flow diagram of an ingress method into the MPLS network 132 is shown. An LSP 144 may initially be established through the MPLS network 132 using conventional signaling protocols (e.g., block 145). During or after establishing the LSP 144, a frame may be received at an ingress edge router 136 (e.g., block 146). The ingress edge router 136 may create the MPLS protocol identification field 102 and the MPLS stack field 104 in the frame while preserving the original layer 3 protocol identification of the frame (e.g., block 148). The ingress edge router 136 may then push one or more MPLS labels 112 onto the MPLS label stack 104 for the LSP 144 (e.g., block 150). The frame in MPLS format may then be forwarded into the MPLS network 132 per the MPLS labels 112 for transmission.

Referring to FIG. 7, a flow diagram of an egress method from the MPLS network 132 is shown. The frame in MPLS format may be transmitted through the MPLS network 132 (e.g., block 154) for reception by an egress edge router 136 (e.g., block 156). The egress edge router 136 may remove the MPLS protocol identification field 102 and the MPLS label stack 104 from the frame (e.g., block 158). The egress edge router 136 is generally not required to reconstruct the original layer 3 protocol identification since the

original layer 3 protocol identification was retained by the ingress edge router 136 (e.g., block 148 of FIG. 6). The egress edge router 136 may then present the frame external to the MPLS network 132 (e.g., block 160).

5 Referring to FIG. 8, a block diagram of an example circuit 162 implementing the ingress and egress method is shown. The circuit 162 generally comprises a computer circuit 164 and a storage medium 166. The computer circuit 164 may connect to the port 137 interfacing to an MPLS network. The computer circuit 164 may connect to the port 138 interfacing to a customer for receiving and sending frames of data. The storage medium 166 may contain a software program 168 defining the ingress method and/or the egress method described above in FIGS. 6 and 7. The software program 168 may be readable and executable by the computer circuit 164 to
15 implement the ingress method and egress method according to the teachings of the present specification, as will be apparent to those skilled in the relevant art(s).

The present invention thus may also include a computer product which may be a storage medium including instructions which
20 can be used to program a computer to perform a process in accordance with the present invention. The storage medium can

0325.00526
CD01223

include, but is not limited to, any type of disk including floppy disk, optical disk, CD-ROM, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, Flash memory, magnetic or optical cards, or any type of media suitable for storing electronic instructions.

5 The present invention may also be implemented by the preparation of ASICs, FPGAs, or by interconnecting an appropriate network of conventional component circuits, as is described herein, modifications of which will be readily apparent to those skilled in the art(s).

10 As used herein, the term "simultaneously" is meant to describe events that share some common time period but the term is not meant to be limited to events that begin at the same point in time, end at the same point in time, or have the same duration.

15 While the invention has been particularly shown and described with reference to the preferred embodiments thereof, it will be understood by those skilled in the art that various changes in form and details may be made without departing from the spirit and scope of the invention.

20 AppleTalk® is a registered trademark of Apple Computers, Inc.